

5. PROTOCOL EXPERIMENT ON TCP-LFN OVER FRAME RELAY

With the growth of the Internet, the TCP/IP protocol has become the most widely used protocol in use today. TCP/IP implementations are commonly found on almost every hardware platform/operating system with a wide variety of applications running over it. Originally designed in the 1960's, the protocol has evolved through the years to meet requirements of LAN's, WAN's, and other systems and networks.

When TCP is used over a satellite link, however, the large bandwidth-delay product can cause problems with throughput. TCP-LFN, an enhanced version of TCP for "long-fat" networks, attempts to rectify these problems [5]. The large bandwidth-delay product of a satellite channel requires larger window sizes to "keep the pipe full" in order to make full use of the channel capacity. Assuming the round-trip time over a satellite link is 0.6 s, the window size required to completely utilize a T1 link would be $1,536,000 \text{ b/s} (0.6 \text{ s} / 8 \text{ b/B}) = 116 \text{ kB}$. Most current TCP implementations have 64 kB for their maximum window size. TCP-LFN increases this limit to $2^{31}-1 \text{ B}$.

TCP retransmits packets when a retransmission timer expires. If this timer expires too soon, packets are unnecessarily retransmitted; if it expires too late, the pipe becomes empty during the intermediate period. Either way, this leads to wasted bandwidth. Current TCP implementations measure one round-trip time (RTT) per window in order to set the retransmission timer. As the window grows, the accuracy of this measurement degrades. TCP-LFN attempts to correct this by measuring an RTT per packet by time-stamping the packet. Another problem that arises for long delays is that the packet sequence numbers can wrap around (be reused). This means that a packet that arrives late has a higher probability of being mistaken for one that was transmitted later with the same sequence number, confusing TCP and delivering incorrect data to the application. TCP-LFN use of time-stamps eliminates the need for sequence numbers, providing protection against wrapped sequence (PAWS) numbers.

5.1 Introduction

The use of frame relay as a WAN protocol has grown significantly over the last few years. Frame relay can support several logical connections over the same physical connection. Although X.25 does this, it uses link-by-link recovery. This type of recovery is wasteful because it leads to lower throughput and requires more complex hardware and software at each network node. Also, with X.25, when the delay is large such as on a satellite link, packets are more likely to be lost and then retransmitted. Frame relay, on the other hand, performs end-to-end recovery leading to a simpler protocol, higher throughput and less complexity in each network node.

Older networks used static bandwidth management to allocate bandwidth to users. More recently, however, there has been a significant growth in the use of bandwidth-on-demand capability since that provides more efficient sharing of network resources, given the bursty nature of data. Also, users

observe less delay since the pipe expands as data flow increases. Users can be charged based on usage as opposed to paying a fixed charge decided at subscription time.

TCP-LFN running over frame relay with bandwidth-on-demand capability is a good solution for networks required for emergency operations that use satellite links. Satellite ground stations are easy to deploy during an emergency and the ability to adjust bandwidth to meet requirements at the emergency site is essential. Satellite networks themselves are less prone to failure since there are fewer physical links and it is possible to enhance link quality dynamically (by coding, etc.). Prioritization of messages is fairly simple to implement within a frame relay network, and is supported by most user devices (such as routers). Most of the user equipment is inexpensive and is commonly available from several vendors; this makes redundant sites more feasible to implement.

This experiment evaluated the performance of TCP-LFN over frame relay using the COMSAT - ACTS frame relay access switch (FRACS) over the ACTS network. The frame relay switches are capable of implementing bandwidth-on-demand over ACTS using the ISDN signaling interface on the ACTS Earth station.

5.2 Experiment Methods and Procedures

TCP packets generated by the application running on the Sun[®] workstation were encapsulated in IP packets and sent over the Ethernet to the WAN router, then they were encapsulated in frame relay packets and sent to the FRACS. The FRACS monitored data flow to each destination periodically and allocated bandwidth by obtaining usage information from the ACTS terminal. The packets were then sent to the ACTS terminal through the T1. The reverse process took place at the destination until the packets reached the application. Figure 5.1 shows the setup for the TCP-LFN over frame relay experiment.

Performance evaluation software tools like TTCP (public domain software) and TCPTGEN (COMSAT proprietary software) were run on the Sun[®] workstations. TTCP opens a TCP socket, sends a specified number of packets of specified size through it, and reports the average throughput measured at the end of the exchange. TCPTGEN is similar except that it measures throughput periodically and also reports the long-term average, therefore it is useful for observing TCP slow start and congestion control mechanisms in action.

The Solaris[®] (Sun[®] operating system) “netstat” command was used to obtain various TCP/IP statistics such as number of data packets sent, maximum segment size, number of retransmissions, retransmission timeout, and fast retransmissions. Another analysis program, SNOOP, was used to observe the following information about packet traces: the times at which each was sent, the outstanding transmit packets, and TCP/IP retransmission behavior.

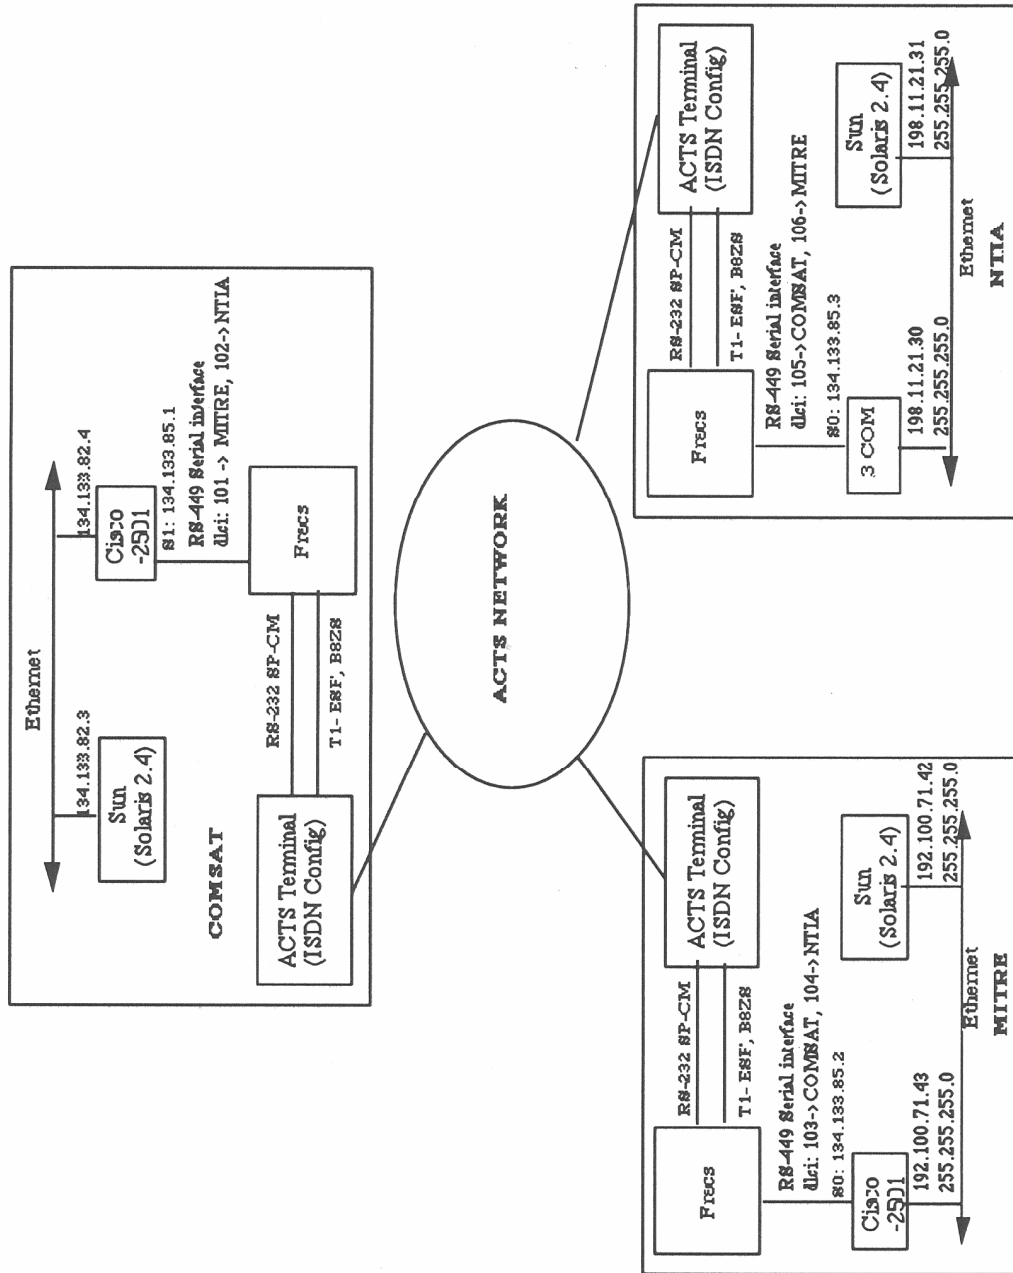


Figure 5.1. Experiment setup.

The FRACS itself also provided a number of valuable statistics, including traffic rate measurements, packet statistics at the data link and trunk levels, link quality measurements, bandwidth and call management statistics, and availability of buffers.

The WAN routers (CISCO and 3-COM) also provided some packet statistics at each of its interfaces, including transmitted and received packets, rate measurements, and buffer information.

5.3 Performance Analysis of TCP-LFN

The performance capabilities of TCP and TCP-LFN were examined using both static and dynamic bandwidth management.

5.3.1 Static Bandwidth Management

Figure 5.2 shows the performance of TCP-LFN when static bandwidth management was used. Both the expected and the measured throughput are as seen by the application. The expected throughput was computed as follows for an application packet size of 1400 B: TCP adds 32 B of overhead, IP adds 20 B, the router adds 4 B, and the FRACS adds 48 B. Hence, the overhead is 104 B or 7%, and the expected throughput is 93% of the allocated bandwidth.

As can be seen in Figure 5.2, the measured throughput is fairly close to the expected throughput. The difference could be made even smaller by using TCP header compression and upgrading the FRACS to use a packet size of 1600 B.

5.3.2 TCP-LFN vs. TCP

Figure 5.3 shows the relative performance of TCP and TCP-LFN with static bandwidth management. The TCP default parameters curve shows the performance of TCP with the standard window size of 8 kB and a maximum (congestion) window size of 32 kB. The throughput is limited to 90 kb/s. Some TCP implementations (e.g., Solaris) allow the user to change some of the TCP driver parameters from their default values. When the transmit and receive window sizes were changed to 64 kB, and the maximum (congestion) window size was changed to 64 kB, the peak throughput increased to 577 kb/s. With TCP-LFN, however, the throughput increased in proportion to the allocated bandwidth.




Figure 5.2. TCP-LFN throughput for the static bandwidth management case.




Figure 5.3. Regular TCP vs. TCP-LFN throughput.

Note that several operating systems do not allow the user to change driver parameters from their default values and those that do often have limits lower than 64 kB (e.g., 50 kB in Solaris[®]) due to implementation/buffer constraints. Also, the application may have to be changed (e.g., “set socket” option in Solaris[®] when opening a TCP socket) or the operating system may have to be rebuilt to do this. Even if this were possible, the larger window sizes would be applied to every connection, not only to those being routed over the satellite. This would result in wasted memory at the hosts; it would also use a large number of buffers at the router, which would lead to longer delays for other connections through the router and less tolerance to short periods of congestion.

5.3.3 Dynamic Bandwidth Management – Preallocated Mode

Figure 5.4 shows the performance of TCP-LFN with dynamic bandwidth management. On the abscissa is the maximum allocated bandwidth. As can be seen in this figure, the measured and the expected TCP-LFN throughput (computed in the same way as in Section 5.3.1) are fairly close. This shows that TCP-LFN works well over a link where bandwidth is allocated on demand. During this measurement, there were no lost packets.




Figure 5.4. TCP-LFN preallocated mode throughput.

In the preallocated mode, a fixed number of channels were reserved from the source to the destination terminal by the FRACS, and these channels were made available to the TCP-LFN traffic as determined by the bandwidth management algorithm. There were several reasons for doing this as opposed to an actual bandwidth-on-demand operation. First, when the ACTS terminal informs the FRACS that a new call has been connected, this may not be true due to latency. Therefore, the first packet(s) sent through this channel may be lost. When this happens, TCP-LFN retransmits the lost packets and reduces the size of its congestion window. As a result the FRACS releases some of the allocated bandwidth. The cycle continues with the throughput rising to about 600 kb/s, and then dropping and rising again. Second, the time taken by the ACTS system to connect a new 64-kb/s call was fairly high (about 5 s). And finally, if there was more than one outstanding call request made to the system it would automatically stop. All of these effects taken together would result in a long delay before the TCP-LFN throughput would rise to its peak value. The preallocated mode alleviates these problems; the results obtained using this mode are the results expected if the above problems did not exist.

5.3.4 Effect of Slow Start

Slow start is a process in TCP-LFN that increases the window size each RTT if no retransmissions are required for the packets sent from the present window. Figure 5.5 shows the time taken by TCP to reach the peak throughput due to slow start (assuming no errors are observed during this period). The peak throughput information was obtained by reviewing the packet traces, finding the congestion window size, and cross-checking it with the periodic measurements done by TCPTGEN.




Figure 5.5. Effect of slow start on throughput.

In this case, several interesting observations were made. First, the round-trip time was 0.8 s, significantly more than the expected value of roughly 0.6 s. The FRACS uses 24 64-kb/s channels (since it may be splitting the T1 amongst several destinations). If a packet is segmented (the mean transmission unit (MTU) on the T1 is 256 B), all segments for the same packet are sent through the same channel. Given a packet size of 1504 B (1400 B of application data + 104 B of headers), the transmission delay is 0.18 s. The rest of the delay (a few milliseconds) is in the transmission delay for the acknowledgment, and the queuing delays in the router, Sun, and FRACS. In order to accommodate this delay, the window size was set to 154,000 B [1,536 kb/s (0.8s / 8 b/B)].

Assuming an initial congestion window and segment size of 1,500 B, a round-trip time of 0.8 s and assuming the congestion window doubles every round-trip time, the peak throughput of 1,536 kb/s (congestion window = 154,000 B) was expected to occur at about 5.6 s (7 round-trip times). As shown in the figure, it actually took 21.18 s for peak throughput to be reached. This is due to the way in which slow start is implemented. Every time the transmitter receives an acknowledgment, the congestion window is incremented by one segment size. This means that if the receiver sent one acknowledgment per data packet, the congestion window would double every round-trip time. However, the receiver sends one acknowledgment for every two (sometimes more) received packets when there is no data flowing in the reverse direction. As a result, the congestion window takes a lot longer to increase to its maximum value. If the transmitter accounted for the number of acknowledged bytes, the congestion window would reach its maximum much earlier. Given that the RTT of a satellite link is significantly larger, slow start severely limits the throughput for transfers that involve a small amount of information.

5.3.5 Effect of Link Errors

Figure 5.6 shows the performance of TCP-LFN with an elevated BER. The errors were injected by putting a bit-error generator on the T1 link between the FRACS and the ACTS terminal.

When the BER was 10^{-6} or worse, the throughput was degraded significantly. When a packet is lost, TCP's fast retransmit algorithm retransmits the packet immediately when a third duplicate acknowledgment is received. TCP then performs congestion avoidance, reducing its congestion window size and slow start threshold (and, hence, the throughput) to half the current value. It then slowly increases its throughput until it reaches the peak value or another error occurs. Using TCPTGEN, this fall and rise of throughput can be observed for each error.

5.4 Results of TCP-LFN over Frame Relay

In this section, the performance of several applications running over TCP-LFN over frame relay, including web browsers, FTP, Telnet, and remote login (Rlogin) is presented.




Figure 5.6. TCP-LFN throughput vs. BER.

5.4.1 Web Browsers

Slow start severely limits the throughput seen by the user for small transfers. Since most browsers (like Netscape[®] and Mosaic[®]) typically open a connection, transfer data from the highlighted link, and then close the connection, the throughput seen by the user is much lower than the link speed. Better performance is seen while transferring large images that are at least several megabytes long.

5.4.2 FTP Performance

Figure 5.7 shows the performance of FTP over TCP-LFN. As can be seen from the figure, the actual throughput at higher bandwidths was lower than that measured using TTCP and TCPTGEN. While the reasons for this were not completely understood, there are several possible reasons. Unlike TCPTGEN, the throughput could not be measured periodically, so the measured throughput was affected by slow start. The size of the file used for the measurements was 9 MB. If larger files had been used, higher throughput may have been observed. The accesses to the hard drive on the computer could also potentially have been a bottleneck at higher throughput.




Figure 5.7. FTP performance over TCP-LFN.

In comparison, when regular TCP was used with the window size set to 50 kB with an allocated bandwidth of 1.536 Mb/s, the measured throughput was 437 kb/s. When multiple FTP sessions were established, the following throughput values were measured.

Table 5.1. Throughput vs. Number of FTP Sessions

Number of FTP Sessions	Throughput (kb/s)
2	776
3	880
4	926

When the number of FTP sessions was increased to five, retransmissions occurred; at six FTP sessions, some packet loss occurred (due to lack of buffers). In both cases, the throughput was about 1 Mb/s.

5.4.3 Telnet and Rlogin

With both Telnet and Rlogin, the longer delay of the satellite link plays a significant role. Operation in the in-line mode can reduce the effects of delay; because this mode produces a local echo of the

characters typed at the terminal. This mode cannot be used for applications such as “vi” hence its usefulness is limited.

5.5 Conclusions

The following conclusions were drawn from the experiment:

The peak throughput that could be attained by TCP over a satellite link was limited to a few hundred kb/s. However, this occurred only when the TCP window sizes were changed to 64 kB. Using the default window size of 8 kB, the peak throughput was less than 100 kb/s. With TCP-LFN, the throughput scaled well with the allocated bandwidth up to 1.544 Mb/s (the limit for this experiment) and can be expected to do so into the gb/s range.

Slow start and the manner in which it is implemented introduces a significant delay in reaching the peak throughput. Most transfers from browsers such as Netscape[®] and Mosaic[®] will not reach peak throughput since they are usually of very short duration and involve setting up and tearing down a connection for each transaction.

TCP and TCP-LFN worked well with bandwidth-on-demand networks with no penalty in the delay or throughput as long as the bandwidth management algorithms were well tuned.

With the introduction of link errors (and no coding on the satellite link), TCP throughput was unaffected at error rates of 10^{-8} or better but degrades significantly at error rates of 10^{-6} or worse.

Interactive applications such as Telnet, which transfer very small amounts of data, suffer mainly from the longer delay. Some limited solutions (line mode) may help alleviate this problem.

6. ACKNOWLEDGMENTS

The authors of this report wish to acknowledge the efforts of a number of individuals who contributed to the progress and success of the work reported. Although not named as authors on the report, these individuals deserve recognition for their contributions.

First, we wish to thank Douglas Vandermade and Quoc Nguyen of MITRE and Paul Redman of COMSAT for their assistance in developing the ACTS Collaboration during its early stages. They worked with William Kissick and Marjorie Weibel of NTIA/ITS and Wayne McCoy and Mary Ruhl of NIST/CSL to develop the first program plans and assisted with early interactions with the sponsor.

Mike Nissley and Kamal Chahine of MITRE replaced Mr. Vandermade and Mr. Nguyen when the experiment phase began. The authors are greatly indebted to Mr. Nissley for ensuring the continued use of the ACTS terminal located at MITRE in Reston, Virginia, and for other administrative assistance. Mr. Nissley and Mr. Chahine both assisted the ACTS Collaboration with experiment planning. Mr. Chahine assisted NTIA/ITS with voice quality measurements and COMSAT with protocol performance measurements. His ability to operate the ACTS Earth station and other specialized, terminal equipment is recognized and appreciated.

Keith Junker of NTIA/ITS prepared the hardware and software for the Voice Quality Assessment System (VQAS). He provided training in its use and ensured that this capability continued to work as expected. He also assisted Dave Sutherland with data collection and reduction.

James Fox of NIST/CSL assisted the NIST authors with the desktop conferencing and LAN bridging experiment planning and execution. He assisted both NIST and NTIA authors with data collection and reduction.

7. REFERENCES

- [1] P. King and S. Voran, "Final report on the NCS voice communications interoperability standards project," Project Report, Sept. 1993.
- [2] J.L. Devore, *Probability and Statistics for Engineering and the Sciences*, Monterey, CA, Brooks/Cole, 1982.
- [3] International Telecommunications Union – Telecommunications Standardization Sector (ITU-T) Recommendation H.320, "Narrow-band visual telephone systems and terminal equipment," Mar. 1993.
- [4] American National Standards Institute (ANSI), American National Standard for Telecommunications – Digital Transport of Video Teleconferencing/Video Telephony Signals, "Video test scenes for subjective and objective performance assessment," Nov. 1995.
- [5] V. Jacobson and R. Braden, "TCP extensions for long-delay paths," RFC 1072, Internet Engineering Task Force, Oct. 1988.

BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION NO. 98-354		2. Gov't Accession No.	3. Recipient's Accession No.
4. TITLE AND SUBTITLE National Security and Emergency Preparedness Communications Experiments Using the Advanced Communications Technology Satellite		5. Publication Date August 1998	
		6. Performing Organization Code NTIA/ITS.N4	
7. AUTHOR(S) William Kissick, et al.		9. Project/Task/Work Unit No.	
8. PERFORMING ORGANIZATION NAME AND ADDRESS National Telecommunications and Information Admin. Institute for Telecommunication Sciences 325 Broadway Boulder, CO 80303-3328		10. Contract/Grant No.	
		12. Type of Report and Period Covered	
11. Sponsoring Organization Name and Address National Communication System (NCS) Office of Programs 701 South Court House Road Arlington, VA 22204-2198		13.	
		14. SUPPLEMENTARY NOTES	
15. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.) Many government telecommunications needs, especially those that support National Security and Emergency Preparedness (NS/EP) missions, are becoming increasingly dependent on commercially available equipment and services. This is consistent with the goals and concepts of the National Information Infrastructure. This report examines the use of an advanced satellite--in this case, NASA's Advanced Communications Technology Satellite (ACTS)--with ISDN and frame relay protocols to support NS/EP communications requirements. A network using three ACTS earth stations was established as a research facility. With this small network, several experiments were performed. Using new objective methods, voice quality was measured over the satellite and compared to other connections such as commercial, terrestrial lines. The performance of applications--desktop conferencing, file transfer, and LAN bridging--that are likely to be useful in NS/EP situations, was determined. The performance of TCP/IP running over frame relay was examined. The results indicate that advanced satellites can be very useful for emergency communications due to the rapidity that earth stations can be deployed, the ease of reconfiguring the satellite, and the practicality of using commonly available applications running over commonly used protocols. However, there are some limitations to the performance of some applications or parts of applications due to the propagation delay of a satellite channel. Telecommunications protocols such as TCP/IP must be significantly modified to perform well over a satellite channel and to take full advantage of bandwidth-on-demand capabilities of an advanced satellite. Key words: Advanced Communications Technology Satellite (ACTS); Integrated Services Digital Network (ISDN); Frame relay; National Security and Emergency Preparedness (NS/EP).			
17. AVAILABILITY STATEMENT <input checked="" type="checkbox"/> UNLIMITED. <input type="checkbox"/> FOR OFFICIAL DISTRIBUTION.		18. Security Class. (This report) Unclassified	20. Number of pages 84
		19. Security Class. (This page) Unclassified	21. Price: